

UK Patent Application (12) GB (19) 2 339 938 (13) A

(43) Date of A Publication 09.02.2000

(21) Application No 9909024.3	(51) INT CL ⁷ G06F 17/30
(22) Date of Filing 21.04.1999	(52) UK CL (Edition R) G4A AUDB
(30) Priority Data (31) 09084452 (32) 26.05.1998 (33) US	(56) Documents Cited WO 98/31155 A2 US 5892917 A US 5732218 A
(71) Applicant(s) International Business Machines Corporation (Incorporated in USA - New York) Armonk, New York 10504, United States of America	(58) Field of Search UK CL (Edition Q) G4A AUDB AUXF AUXX INT CL ⁶ G06F 11/34 17/30 17/40 Online: WPI, EPODOC, INSPEC, COMPUTER
(72) Inventor(s) Gennaro A Cuomo Binh Q Nguyen Sandeep K Singhal	
(74) Agent and/or Address for Service IBM United Kingdom Limited Intellectual Property Department, Mail Point 110, Hursley Park, WINCHESTER, Hampshire, SO21 2JN, United Kingdom	

(54) Abstract Title

Collecting information about document retrievals over the World Wide Web

(57) Collecting information about which users are accessing different sites on the World Wide Web, and what site content those users are accessing, e.g. to enable "profiles" to be established for each site visitor and tuning of the Web site content to meet the visitors' interests. When a Web server supports server-side scripting (Fig. 2), wherein the script receives a Universal Resource Identifier (URI) and the requesting user identity to dynamically generate the document content that should be returned to the user, a particular URI is not associated with unique Web site content and so cannot be used alone for developing meaningful user profiles. Instead, according to the invention, a document key is also logged, wherein retrieved documents which are substantially similar (i.e. with only minor variations in their content, such as in the requesting user's name) are assigned the same document key. A wide range of metrics may be used in order to compare two documents for similarity. Different dynamically-generated documents retrieved using the same URI are thus distinguished, while also merging access information about documents that are nearly identical.




FIG. 4

GB 2 339 938 A




FIG. 2




FIG. 3





FIG. 4



METHOD AND SYSTEM FOR WORLD-WIDE WEB INFORMATION COLLECTION

5 The present invention relates to the collection of information, and in particular to collecting information about document retrievals over the World-Wide Web.

10 In the World-Wide Web, a content provider deploys a plurality of Web servers that deliver Web pages to clients. When requesting a Web page, the client supplies a Uniform Resource Locator (URL) or Universal Resource Identifier (URI) to the server. The server associates this URI with a particular page of content and delivers that information to the requesting client.

15 As the World-Wide Web is being used increasingly to support commerce and targeted advertising, content providers desire to collect information about which users are accessing the site and what site content those users are accessing. This information can be used to establish "profiles" for each site visitor and enable tuning of the Web site content to meet the visitors' interests. Traditionally, this visitor information is collected by the Web server or a proxy server in the form of a log file. This log file contains, among other things, the requesting host address, the requested URI, and the time at which the request was received. Because each URI represents a particular piece of static content at the Web site, the URI is sufficient for a user profile analyzer to evaluate which content was received by each user and to detect similarities among the behavior of different users.

20 Recent Web servers are providing support for server-side scripting, whereby the URI is associated with a program or script that is executed at the Web server. This script is responsible for receiving the URI and the user identity and using this information to dynamically generate the content that should be returned to the requesting user. This generated content may account for the user's previous behavior at the site, his/her access permissions, his/her demographic information, or any number of other factors. Dynamic server content is supported by most Web servers today, including Microsoft's Active Server Pages, Sun's Dynamic Server Pages, industry-standard servlets, Common Gateway Interface (CGI) executables, and other mechanisms.

25 As a result of this direction, a particular URI can no longer be associated with particular content at the Web site. On different requests, the URI may return wholly different content depending on the requesting user and the context in which the request was issued. Consequently, existing methods for capturing user information are insufficient for producing meaningful user profiles. More specifically,

the reliance on URIs alone prevents the accurate characterization of which users are exhibiting similar access behavior.

5 Accordingly, the invention provides a method of collecting information about document retrievals over the World-Wide Web, comprising the steps of:

receiving a requesting user identity, requested Universal Resource Identifier (URI), and a content of a retrieved document;

10 selecting a Candidate Document from a Retrieved Document Database, said Candidate Document associated with a Candidate Document Key;

comparing said retrieved document to said Candidate Document to determine their similarity;

associating said retrieved document with a Retrieved Document Key; and

15 adding a Log File Entry including said requesting user identity, said requested URI, and said Retrieved Document Key.

20 In the preferred embodiment, said Retrieved Document Key is equal to said Candidate Document Key if said Candidate Document is deemed to be similar to the retrieved document, and said Retrieved Document Key is newly generated if said Candidate Document is deemed to be dissimilar to the retrieved document. The step of associating said retrieved document with a Retrieved Document Key further comprises the step of adding said retrieved document to said Received Document Database. Each of a 25 plurality of documents in said Retrieved Document Database is associated with a Document Comparator and a first Document Comparator may be compared to a second Document Comparator using a Document Comparator Function.

30 Typically this comparison is performed by the steps of: computing said first Document Comparator for said retrieved document; retrieving said second Document Comparator for said Candidate Document; computing with said Document Comparator Function a numeric measure of the difference between said first Document Comparator and said second Document Comparator; and comparing said numeric measure against a 35 predefined Document Difference Threshold.

40 There are various possibilities for forming the Document Comparator. For example, it may comprise content or a list of significant words or phrases from the document associated therewith. Alternatively, the Document Comparator may be computed by associating predefined portions of the document associated therewith to a binary token. Each said Document Comparator may comprise a Comparator for each of a plurality of predefined sections of the document associated therewith.

5 In the preferred embodiment, said step of selecting a Candidate Document comprises selecting from a Document Comparator Database, and the URI for said Candidate Document is equal to the URI for said retrieved document. Said steps of selecting a Candidate Document and comparing said retrieved document to said Candidate Document are repeated until a Candidate Document is determined to be similar to the retrieved document, or no further Candidate Documents are available.

10 The invention further provides a system for collecting information about document retrievals over the World-Wide Web, comprising:

 means for receiving a requesting user identity, requested Universal Resource Identifier (URI), and a content of a retrieved document;

 means for selecting a Candidate Document from a Retrieved Document Database, said Candidate Document associated with a Candidate Document

15 Key;

 means for comparing said retrieved document to said Candidate Document to determine their similarity;

 means for associating said retrieved document with a Retrieved Document Key; and

20 means for adding a Log File Entry including said requesting user identity, said requested URI, and said Retrieved Document Key.

25 The invention further provides a computer program product recorded on a computer readable medium for collecting information about document retrievals over the World-Wide Web, comprising:

 means for receiving a requesting user identity, requested Universal Resource Identifier (URI), and a content of a retrieved document;

 means for selecting a Candidate Document from a Retrieved Document Database, said Candidate Document associated with a Candidate Document Key;

 means for comparing said retrieved document to said Candidate Document to determine their similarity;

 means for associating said retrieved document with a Retrieved Document Key; and

30 means for adding a Log File Entry including said requesting user identity, said requested URI, and said Retrieved Document Key.

40 It will be appreciated that such a system and computer program product can usefully employ the same set of preferred features as described in relation to the method of the invention.

45 The approach described herein allows the efficient collection of user access information in the presence of dynamically-generated content at a Web server, in order to support the accurate generation of user profiles. In particular, profile information about users accessing Web

pages from a plurality of Web servers can be collected when the Web content is generated dynamically for each request at the Web server.

5 Thus each user's request within a networked environment for World-Wide Web information can be associated with the content of the retrieved document when that document was generated dynamically by analyzing the content of retrieved documents and associating Document Comparators with each document. This approach allows user requests that retrieve the same document content to be grouped together, ignoring minor 10 variations in document content as might occur when the documents differ only in the presence of the requesting user's name. A wide range of metrics may be used for comparing two documents for similarity.

15 A preferred embodiment will now be described in detail by way of example only with reference to the following drawings:

Figure 1 is a pictorial representation of a data processing system;

Figure 2 shows a block diagram of a World-Wide Web environment in which user access information may be collected in accordance with the present invention;

20 Figure 3 shows a sample data structure for representing the information collected in accordance with the present invention; and

Figure 4 is a flowchart showing how an Access Information Collector analyzes a document retrieved from a Web server and updates its data structures.

25 Referring to Figure 1, there is depicted a graphical representation of a data processing system 8. As may be seen, data processing system 8 may include a plurality of networks, such as Local Area Networks (LAN) 10 and 32, each of which preferably includes a plurality of individual computers 12 and 30, respectively. Of course, those skilled in the art will appreciate that a plurality of Intelligent Work Stations (IWS) coupled to a host processor may be utilized for each such network. Each said network may also consist of a plurality of processors coupled via a communications medium, such as shared memory, shared storage, or an interconnection network. As is common in such data processing systems, each individual computer may be coupled to a storage device 14 and/or a printer/output device 16 and may be provided with a pointing device such as a mouse 17.

40 The data processing system 8 may also include multiple mainframe computers, such as mainframe computer 18, which may be preferably coupled to LAN 10 by means of communications link 22. The mainframe computer 18 may also be coupled to a storage device 20 which may serve as remote storage for LAN 10. Similarly, LAN 10 may be coupled via communications link 24 through a sub-system control unit/communications controller 26

and communications link 34 to a gateway server 28. The gateway server 28 is preferably an IWS which serves to link LAN 32 to LAN 10.

With respect to LAN 32 and LAN 10, a plurality of documents or resource objects may be stored within storage device 20 and controlled by mainframe computer 18, as resource manager or library service for the resource objects thus stored. Of course, those skilled in the art will appreciate that mainframe computer 18 may be located a great geographic distance from LAN 10 and similarly, LAN 10 may be located a substantial distance from LAN 32. For example, LAN 32 may be located in California while LAN 10 may be located within North Carolina and mainframe computer 18 may be located in New York.

Software program code is typically stored in the memory of a storage device 14 of a stand alone workstation or LAN server from which a developer may access the code for distribution purposes. The software program code may be embodied on any of a variety of known media for use with a data processing system such as a diskette or CD-ROM or may be distributed to users from a memory of one computer system over a network of some type to other computer systems for use by users of such other systems. Such techniques and methods for embodying software code on media and/or distributing software code are well-known and will not be further discussed herein.

Referring now to Figure 2, components of a World-Wide Web system are shown in which user information may be gathered in accordance with the present invention. A plurality of clients (generally indicated by reference numerals 200, 201, and 202) access information over a network 205 using World-Wide Web browsers such as NETSCAPE NAVIGATOR, a trademark of Netscape Communications Corporation, or MICROSOFT INTERNET EXPLORER, a trademark of Microsoft Corporation. These clients access a plurality of Web servers (generally indicated by reference numerals 210, 211, and 212) such as LOTUS GO, a trademark of Lotus Corporation, MICROSOFT INTERNET INFORMATION SERVICE (IIS), a trademark of Microsoft Corporation, or NETSCAPE FASTTRACK, a trademark of Netscape Communication Corporation. In accessing these Web servers, the clients 200, 201 and 202 specify a URI. Each of these Web servers 210, 211, and 212 accesses a Static Content Database (generally indicated by reference numerals 220, 221, and 222) and a Dynamic Content Generator (generally indicated by reference numerals 230, 231, and 232) that receives a URI and other information about the user and generates Web content suitable for display by the browsers at the clients 200, 201, and 202. These Dynamic Content Generators 230, 231, and 232 may take many forms, including Active Server Pages, servlets, Common Gateway Interface (CGI) binaries, or Dynamic Server Pages.

Upon receiving a URI request from a client, the Web server 210, 211, or 212 either retrieves the content from the Static Content Database 220, 221, or 222 or from the Dynamic Content Generator 230, 231, or 232. An Access Information Collector 240 receives client requests and content returned from the Static Content Database 220, 221, or 222 or from the Dynamic Content Generator 230, 231, or 232 and collects log information that can be used to analyze the access patterns of various users.

It should be understood that the physical location of the components shown in Figure 2 may vary. In particular, the Access Information Collector 240 may be embedded in the Web servers 210, 211, and 212. Moreover, the Dynamic Content Generators 200, 201, and 202 and Static Content Databases 220, 221, and 222 may be co-located with the Web servers 210, 211, and 212.

Figure 3 illustrates the information collected by the Access Information Collector. A Log File 300 contains a sequence of Access Records. Each Access Record includes at least a time stamp 301, a requested URI 313, and a Document Key 312.

A Retrieved Document Database 310 contains a repository of Document Records corresponding to documents retrieved by users. Each Document Record 311 is indexed by a Document Key 312 and contains an associated URI 313, document text 314, and a Document Comparator 315. The Document Key 312, when combined with the URI 313, serves to uniquely identify the Document Record 311. Document Keys may be assigned sequentially or by any other appropriate method.

The Document Comparator 315 is a representation of the document's contents and is used by a Document Comparator Function to determine whether there are substantial predefined similarities, as will be subsequently described in greater detail, between the current document and other previously retrieved documents. The Document Comparator Function receives the Document Comparators for two documents and determines whether the two documents are substantially similar. To make this determination, the Function may employ a Document Difference Threshold, a numeric value that indicates how much two documents may differ before they are no longer deemed to be substantially similar. The use of the Document Difference Threshold depends on the particular Document Comparator Function being used. The use of a Document Difference Threshold allows the Document Comparator Function to ignore minor differences between two documents. Such minor differences include timestamps, client name, or client-specific data.

In the present embodiment of this invention, the Document Comparator 315 is the actual content of the document itself, and the

Document Comparator Function for any two documents is defined to be the number of character insertions, deletions, or modifications required to convert one document to the other. This computation is well understood in the prior art (see, for example, the use of tries, as described in 5 Chapter 11 of Alan Tharp, *File Organization and Processing*, Wiley, 1988) and will not be discussed further. Alternative embodiments of this invention may compute a Document Comparator 315 by mapping each word, paragraph, or section of the document to a binary token. In this case, the Document Comparator Function might count the number of matching 10 binary tokens, and the Document Difference Threshold would designate what percentage of the tokens must match (see, for example, "Copy Detection Mechanisms for Digital Documents," by Sergey Brin, James Davis, and Hector Garcia-Molina, in *Proceedings of the 1995 SIGMOD International Conference on Management of Data*, pages 398-409, May 1995). Yet another 15 embodiment of this invention may define a Document Comparator 315 as a list of the most significant (as predefined) words or phrases in the document; the Document Comparator Function may simply count how many words or phrases occur in both documents, and the Document Difference Threshold would designate what percentage of words in each document must appear in the other. Other comparison methods are well established in 20 the prior art. The essential element of a Document Comparator 315 is that a metric (i.e. the Document Comparator Function) must exist for comparing two different Document Comparators to determine by how much 25 their respective documents differ. Indeed, a Document Comparator 315 may actually comprise multiple Comparators, one per each predefined section of the document, each having an associated Document Comparator Function.

Finally, a Document Comparator Index 320 associates each Document Comparator 315 with the corresponding Document Key 312. The Index 320 is 30 used to improve the performance of the Document Comparator 315 evaluations and the selection of Candidate Documents (see Figure 4). However, it is a performance optimization that may be omitted by 35 alternative embodiments of this invention.

Though the data structures have been illustrated in Figure 3 with a 35 particular embodiment, alternative representations of this information are possible. The main attributes used are the association of each Document Comparator 315 to a Document Key 312, the association of each user URI 313 retrieval with a particular Document Key 312, and the 40 association of each Document Key 312 with particular document content. It should be noted that various optimizations are also possible. For example, instead of storing each document's full content, the Retrieved Document Database 310 may store only a list of most significant words or 45 phrases.

When a document is accessed from the Web server (with a particular URI), the Access Information Collector 240 analyzes the retrieved document (using the Document Comparator Function) to determine whether it is substantially similar to another document that has been previously retrieved from that Web server using the same URI. If a substantially similar document has already been generated by the Web server, then the user's access is associated with that previous document; however, if a substantially similar document has not been previously generated by the Web server, then the user's access is associated with this new document. In this way, the Access Information Collector 240 distinguishes between different dynamically-generated documents retrieved using the same URI while also merging access information about documents that are nearly identical.

Referring now to Figure 4, a flowchart depicts the steps taken by the Access Information Collector 240 to analyze a document retrieved from a Web server and to update the Log File 300, Retrieved Document Database 310, and Document Comparator Index 320 (as shown in Figure 3). At block 400, the Access Information Collector 240 receives the requested URI, the time of the request, the identity of the requesting client, and the content of the retrieved document. At block 402, a Document Comparator 315 is computed for the retrieved document. At block 404, a Candidate Document and Candidate Document Comparator are selected from the Retrieved Document Database 310. The Candidate Document is a document in the Retrieved Document Database 310 whose URI matches that of the retrieved document. (It should be understood that alternative embodiments may remove the restriction that the URI of the retrieved document and the URI of the Candidate Document match, or may introduce additional restrictions on what constitutes a Candidate Document). At decision block 406, it is determined whether or not a Candidate Document has been found. If the answer to decision block 406 is yes, then at decision block 408, the Document Comparator Function is invoked with the Document Comparators of the retrieved document and of the Candidate Document to determine whether or not the retrieved document and the Candidate Document are substantially similar.

Continuing with Figure 4, if the answer to decision block 408 is yes, then it is determined that the retrieved document is sufficiently similar to the Candidate Document and no new entry is required to either the Retrieved Document Database 310 or to the Document Comparator Index 320. At block 410, the Document Key is retrieved for the Candidate Document. At block 415, a new entry is added to the Log File, including the time stamp, requested URI, and candidate document's Document Key. The process then terminates at block 490. If the answer to decision block 408 is no, then control returns to block 404, where another Candidate Document is selected for evaluation.

5 If the answer to decision block 406 is no, then it is determined that the retrieved document is new. At block 420, a new Document Key is generated for the retrieved document. At block 425, a new entry is added to the Retrieved Document Database 310 to associate the retrieved document's Document Key with a new Document Record containing the retrieved URI, retrieved document, and retrieved document's Document Comparator. At block 430, a new entry is added to the Document Comparator Index 320 database to associate the retrieved document's Document Comparator with the retrieved document's Document Key. At block 10 435, a new entry is added to the Log File, including the time stamp, requested URI, and retrieved document's Document Key. The process then terminates at block 490.

15 Thus, each user access is associated with a Document Key representing a document in the Retrieved Document Database with a sufficiently close Document Comparator. Each URI is, therefore, potentially linked with multiple documents, each having different content. At the same time, the analysis ignores minor differences between documents, as might arise when page content is customized in minor ways to reflect the identity of the requesting user.

20

CLAIMS

1. A method of collecting information about document retrievals over the World-Wide Web, comprising the steps of:

5 receiving a requesting user identity, requested Universal Resource Identifier (URI), and a content of a retrieved document;

selecting a Candidate Document from a Retrieved Document Database, said Candidate Document associated with a Candidate Document Key;

comparing said retrieved document to said Candidate Document to determine their similarity;

10 associating said retrieved document with a Retrieved Document Key; and

adding a Log File Entry including said requesting user identity, said requested URI, and said Retrieved Document Key.

15 2. The method of Claim 1, wherein said Retrieved Document Key is equal to said Candidate Document Key if said Candidate Document is deemed to be similar to the retrieved document.

20 3. The method of Claim 1 or 2, wherein said Retrieved Document Key is newly generated if said Candidate Document is deemed to be dissimilar to the retrieved document, and wherein said step of associating said retrieved document with a Retrieved Document Key further comprises the step of adding said retrieved document to said Received Document Database.

25 4. The method of any preceding Claim, wherein each of a plurality of documents in said Retrieved Document Database is associated with a Document Comparator and wherein a first Document Comparator may be compared to a second Document Comparator using a Document Comparator Function.

30 5. The method of Claim 4, wherein said step of comparing to determine the similarity of said Candidate Document to the retrieved document further comprises the steps of:

35 computing said first Document Comparator for said retrieved document;

retrieving said second Document Comparator for said Candidate Document;

40 computing with said Document Comparator Function a numeric measure of the difference between said first Document Comparator and said second Document Comparator; and

45 comparing said numeric measure against a predefined Document Difference Threshold.

6. The method of Claim 4 or 5, wherein each said Document Comparator comprises content from the document associated therewith.

5 7. The method of Claim 4 or 5, wherein each said Document Comparator is computed by associating predefined portions of the document associated therewith to a binary token.

10 8. The method of Claim 4 or 5, wherein each said Document Comparator comprises a list of significant words or phrases from the document associated therewith.

15 9. The method of any of Claims 4-8, wherein each said Document Comparator comprises a Comparator for each of a plurality of predefined sections of the document associated therewith.

20 10. The method of any of Claims 4-9, wherein said step of selecting a Candidate Document comprises selecting from a Document Comparator Database.

25 11. The method of any preceding Claim, wherein the URI for said Candidate Document is equal to the URI for said retrieved document.

30 12. The method of any preceding Claim, wherein said steps of selecting a Candidate Document and comparing said retrieved document to said Candidate Document are repeated until a Candidate Document is determined to be similar to the retrieved document, or no further Candidate Documents are available.

35 13. A system for collecting information about document retrievals over the World-Wide Web, comprising:
means for receiving a requesting user identity, requested Universal Resource Identifier (URI), and a content of a retrieved document;
means for selecting a Candidate Document from a Retrieved Document Database, said Candidate Document associated with a Candidate Document Key;
means for comparing said retrieved document to said Candidate Document to determine their similarity;
means for associating said retrieved document with a Retrieved Document Key; and
40 means for adding a Log File Entry including said requesting user identity, said requested URI, and said Retrieved Document Key.

45 14. The system of Claim 13, wherein said Retrieved Document Key is equal to said Candidate Document Key if said Candidate Document is deemed to be similar to the retrieved document.

15. The system of Claim 13 or 14, wherein said Retrieved Document Key is newly generated if said Candidate Document is deemed to be dissimilar to the retrieved document, and wherein said means for associating said retrieved document with a Retrieved Document Key further comprises means for adding said retrieved document to said Received Document Database.

5

16. The system of any of Claims 13-15, wherein each of a plurality of documents in said Retrieved Document Database is associated with a Document Comparator and wherein a first Document Comparator may be compared to a second Document Comparator using a Document Comparator Function.

10

17. The system of Claim 16, wherein said means for comparing to determine the similarity of said Candidate Document to the retrieved document further comprises:

15

means for computing said first Document Comparator for said retrieved document;

means for retrieving said second Document Comparator for said Candidate Document;

20

means for computing with said Document Comparator Function a numeric measure of the difference between said first Document Comparator and said second Document Comparator; and

means for comparing said numeric measure against a predefined Document Difference Threshold.

25

18. The system of Claim 16 or 17, wherein each said Document Comparator comprises content from the document associated therewith.

30

19. The system of Claim 16 or 17, wherein each said Document Comparator is computed by associating predefined portions of the document associated therewith to a binary token.

35

20. The system of Claim 16 or 17, wherein each said Document Comparator comprises a list of significant words or phrases from the document associated therewith.

40

21. The system of any of Claims 16-20, wherein each said Document Comparator comprises a Comparator for each of a plurality of predefined sections of the document associated therewith.

45

22. The system of any of Claims 16-21, wherein said means for selecting a Candidate Document comprises means for selecting from a Document Comparator Database.

23. The system of any of Claims 13-22, wherein the URI for said Candidate Document is equal to the URI for said retrieved document.

24. The system of any of Claims 13-23, wherein said means for selecting a Candidate Document and means for comparing said retrieved document to said Candidate Document operate repetitively until a Candidate Document is determined to be similar to the retrieved document, or no further Candidate Documents are available.

5

25. A computer program product recorded on a computer readable medium for collecting information about document retrievals over the World-Wide Web, comprising:

10 means for receiving a requesting user identity, requested Universal Resource Identifier (URI), and a content of a retrieved document;

means for selecting a Candidate Document from a Retrieved Document Database, said Candidate Document associated with a Candidate Document Key;

15 means for comparing said retrieved document to said Candidate Document to determine their similarity;

means for associating said retrieved document with a Retrieved Document Key; and

20 means for adding a Log File Entry including said requesting user identity, said requested URI, and said Retrieved Document Key.

26. The product of Claim 25, wherein said Retrieved Document Key is equal to said Candidate Document Key if said Candidate Document is deemed to be similar to the retrieved document.

25

27. The product of Claim 25 or 26, wherein said Retrieved Document Key is newly generated if said Candidate Document is deemed to be dissimilar to the retrieved document, and wherein said means for associating said retrieved document with a Retrieved Document Key further comprises means for adding said retrieved document to said Received Document Database.

30

28. The product of any of Claims 25-27, wherein each of a plurality of documents in said Retrieved Document Database is associated with a Document Comparator and wherein a first Document Comparator may be compared to a second Document Comparator using a Document Comparator Function.

35

29. The product of Claim 28, wherein said means for comparing to determine the similarity of said Candidate Document to the retrieved document further comprises:

40 means for computing said first Document Comparator for said retrieved document;

means for retrieving said second Document Comparator for said Candidate Document;

means for computing with said Document Comparator Function a numeric measure of the difference between said first Document Comparator and said second Document Comparator; and

5 means for comparing said numeric measure against a predefined Document Difference Threshold.

30. The product of Claim 28 or 29, wherein each said Document Comparator comprises content from the document associated therewith.

10 31. The product of Claim 28 or 29, wherein each said Document Comparator is computed by associating predefined portions of the document associated therewith to a binary token.

15 32. The product of Claim 28 or 29, wherein each said Document Comparator comprises a list of significant words or phrases from the document associated therewith.

20 33. The product of any of Claims 28-32, wherein each said Document Comparator comprises a Comparator for each of a plurality of predefined sections of the document associated therewith.

34. The product of any of Claims 28-33, wherein said means for selecting a Candidate Document comprises means for selecting from a Document Comparator Database.

25 35. The product of any of Claims 25-34, wherein the URI for said Candidate Document is equal to the URI for said retrieved document.

30 36. The product of any of Claims 25-35, wherein said means for selecting a Candidate Document and means for comparing said retrieved document to said Candidate Document operate repetitively until a Candidate Document is determined to be similar to the retrieved document, or no further Candidate Documents are available.



Application No: GB 9909024.3
Claims searched: 1-36

15

Examiner: Melanie Gee
Date of search: 23 November 1999

Patents Act 1977
Search Report under Section 17

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:

UK Cl (Ed.Q): G4A (AUDB, AUXF, AUXX)

Int Cl (Ed.6): G06F 11/34, 17/30, 17/40

Other: Online: WPI, EPODOC, INSPEC, COMPUTER

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
A, P	WO 98/31155 A2 (RELEVANT KNOWLEDGE)	
A, P	US 5892917 A (MYERSON)	
A	US 5732218 A (BLAND et al.)	

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

**This Page is Inserted by IFW Indexing and Scanning
Operations and is not part of the Official Record**

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images include but are not limited to the items checked:

- BLACK BORDERS**
- IMAGE CUT OFF AT TOP, BOTTOM OR SIDES**
- FADED TEXT OR DRAWING**
- BLURRED OR ILLEGIBLE TEXT OR DRAWING**
- SKEWED/SLANTED IMAGES**
- COLOR OR BLACK AND WHITE PHOTOGRAPHS**
- GRAY SCALE DOCUMENTS**
- LINES OR MARKS ON ORIGINAL DOCUMENT**
- REFERENCE(S) OR EXHIBIT(S) SUBMITTED ARE POOR QUALITY**
- OTHER:** _____

IMAGES ARE BEST AVAILABLE COPY.

As rescanning these documents will not correct the image problems checked, please do not report these problems to the IFW Image Problem Mailbox.

This Page Blank (uspto)